

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

ANALYSIS OF MESSAGE SEQUENCES

Inventors:

John R. Lambert

and

Luis Felipe Cabrera

ATTORNEY'S DOCKET NO. MS1-1714US

TECHNICAL FIELD

This subject matter relates to automated analysis techniques, and in a more particular implementation, to automated techniques for investigating the behavior of data processing systems, such as computer systems.

BACKGROUND

Analysts commonly apply one or more techniques for investigating the behavior of data processing systems. An analyst may apply such techniques to determine whether a data processing system is working properly. Functional tests ensure that the data processing system is producing expected results. Performance-related tests ensure that the data processing system is producing the expected results in a desired manner (such as within a particular period of time, etc.). Alternatively, the analyst may apply investigation techniques in an open-ended manner to explore the behavior of the data processing system to determine its salient characteristics (e.g., without necessarily comparing this behavior with predefined expectations). These techniques can be applied to any kind of data processing system, included computers running software programs, networks of such computers, data processing equipment included hardwired (non-programmable) processing logic, or other kinds of processing device(s).

An analyst can select from a great variety of strategies in investigating the behavior of a data processing system. Many of these strategies require *a priori* knowledge of the features of the system under investigation and its output. One class of such techniques constructs a model of the system under consideration to provide a baseline that defines the expected behavior of the system. This class of techniques then measures the actual behavior of the system and compares it with the baseline model. Discrepancies between measured and expected results may suggest that the system is not

1 working properly. For instance, such a technique may analyze the messages output from
2 a data processing system under test and then compare such messages with a model that
3 defines the expected form and content of such messages to determine whether the system
4 is operating properly.

5 The above-described solution may not be able to diagnose problems in some
6 kinds of data processing systems. Consider, for example, the case of a data processing
7 system that includes multiple computer devices interacting with each other via a network.
8 Two computers may be transmitting messages with each other that have the correct data
9 type and content. Nevertheless, the timing at which these messages are being transmitted
10 and received, or the ordering or number of such messages, may suggest that there is some
11 anomaly within the data processing system; this anomaly cannot be detected by simply
12 examining the form of each individual message being transmitted. Furthermore, an
13 analyst may wish to investigate the behavior of a data processing system that the analyst
14 cannot gain direct access to, and therefore the analyst may not know the details of its
15 configuration. Therefore, the analyst may be unaware, beforehand, of what messages and
16 message sequences are valid (properly formed) and what messages and message
17 sequences are invalid (improperly formed).

18 Another class of investigation techniques may apply formal methods of message
19 analysis based on a finite state machine. However, it may be difficult or impossible to
20 construct such a state machine for many data processing machines. It may be particularly
21 difficult to construct such a model where the behavior of the system is non-deterministic,
22 or where the model must also account for systems which permit message retries. Further,
23 as in the first class of techniques, building a finite state machine requires advance
24 knowledge of the configuration of the data processing system. This class of techniques
25 therefore does not work in cases where the analyst cannot determine the configuration of

1 the data processing system (because, for instance, the data processing system is a network
2 resource that is owned and maintained by an entity not under the control of the analyst).

3 Another class of techniques captures some kind of code profile of the system
4 under consideration, such as an operational profile or execution profile. These techniques
5 then analyze various features in the profile. For example, one known technique analyzes
6 the behavior of a standalone system by applying test instrumentation to count function
7 calls. This test instrumentation can be implemented with code that interacts with the code
8 of the system under test. There are drawbacks to this class of techniques as well. For
9 instance, this solution requires invasive instrumentation to monitor the internal behavior
10 the data processing system. Again, where the data processing system is not under control
11 of the analyst, this solution might not be possible or feasible. Further, different data
12 processing systems may adopt different versions of a software program. In this case, test
13 instrumentation adapted to interact with one version of the software program might not
14 work well (or at all) with another version of the software program. Further, the test
15 results generated by one version may not be directly comparable to the test results
16 generated by another version of the program. These differences complicate the
17 monitoring and analysis of the behavior of the system, because the analyst must
18 specifically tailor his or her test strategy to account for these differences (such as by
19 selecting test instrumentation that is adapted to work with different versions, and then
20 harmonizing the test results between different versions).

21 As such, there is an exemplary need in the art for a more efficient, effective,
22 and/or flexible technique for investigating the operational characteristics of data
23 processing systems.
24
25

SUMMARY

According to one exemplary implementation, a method is described for investigating messages passed in a message-passing environment. The method can involve: (1) collecting a plurality of messages from at least one participant in the message-passing environment; (2) assembling the messages into at least one message sequence; (3) analyzing said at least one message sequence to extract information regarding the message-passing environment; and (4) outputting the information to a user.

A related apparatus and computer readable media are also described herein.

In some message-passing environments, the messages can be intercepted at locations between participants in the message exchange. Accordingly, this analysis technique may not need to account for the configuration complexities of any participant. Further, in some environments, this analysis technique may work even though the analyst does not have access to the systems used by one or more participants in the message-passing environment. Additional benefits of this approach are identified in the following discussion.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows an exemplary system for investigating the behavior of a data processing environment by analyzing messages passed between participants in this environment.

Fig. 2 shows four exemplary data processing environments that the system of Fig. 1 can be applied to.

Fig. 3 shows exemplary message analysis logic and a message sequence data store for use in the system of Fig. 1.

1 Fig. 4 shows an exemplary method for investigating the behavior of a data
2 processing environment using, for instance, the system of Fig. 1.

3 Fig. 5 shows an exemplary output of the method shown in Fig. 4.

4 Fig. 6 shows an exemplary computing environment for implementing the system
5 of Fig. 1.

6 The same numbers are used throughout the disclosure and figures to reference like
7 components and features. Series 100 numbers refer to features originally found in Fig. 1,
8 series 200 numbers refer to features originally found in Fig. 2, series 300 numbers refer
9 to features originally found in Fig. 3, and so on.

10 11 **DETAILED DESCRIPTION**

12 **A. Exemplary System for Performing Message-Based Analysis**

13 Fig. 1 shows an exemplary system 100 for investigating a message-passing
14 environment 102. By way of overview, the message-passing environment 102 is shown
15 as including at least two participants (104, 106). These participants (104, 106) transmit
16 messages (M) to each other (or, in some cases, to multiple participants in broadcast
17 mode, and in other cases, to themselves). An analysis system 108 collects these
18 messages via various observation agents (O) (e.g., 110, 112, 114, 116) and then groups
19 them into sequences for storage in a data store 118. Message analysis logic 120 analyzes
20 these message sequences and forms an output result based thereon.

21 The output result can provide insight into the behavior of the message-passing
22 environment 102. For instance, the output result may group similar message sequences
23 together using cluster analysis or some other technique. From this cluster analysis, the
24 message analysis logic 120 can provide an indication of any message sequences which
25 may differ substantially from others. These outlying message sequences may represent

1 an anomalous and undesired condition within the message-passing environment 102.
2 More specifically, the anomalous condition may suggest that certain modules of the
3 message-passing environment 102 are outputting incorrect results, or are providing
4 correct results yet providing the results in an inefficient manner (e.g., either by taking too
5 long to provide the results or by consuming too much system resources in generating the
6 results). Corrective action can be taken on the basis of the output of the message analysis
7 logic 120.

8 The above-described analysis strategy has numerous advantages compared to the
9 kinds of techniques described in the Background section of this disclosure. For instance,
10 analysis is based on the flow of messages passed between participants, rather than an in-
11 depth knowledge of the configuration of each participant. Hence, meaningful
12 information can be extracted from the message-passing environment even though the
13 analyst does not know the precise configuration of each participant. Indeed, the analyst
14 might not even have knowledge of the identity of an entity sending or receiving a
15 message (as well as any intermediary agents that may process the message en route from
16 sender to receiver). This aspect of the strategy simplifies the investigation because the
17 analyst need no longer generate a model of the system being tested in order to analyze its
18 behavior. An analyst also need not be concerned when participants are running different
19 versions of a common software product, as the investigation is based on the
20 communication between participants, rather than the configuration of each participant *per*
21 *se.*

22 Further, in some cases, messages can be collected at locations “on the wire”
23 between participants. Thus, an analyst might be able to collect meaningful information
24 from the message-passing environment 102 even though the analyst does not have
25 authority or the ability to directly access the systems provided by each participant. This

1 is a particularly attractive feature when analyzing behavior of wide area network systems
2 based on traffic on the network, as the messages may originate and pass through a great
3 number of processing agents that are not under the direct control of the analyst.

4 The reader will appreciate that there are additional merits to the system and
5 method described herein.

6 After the above overview, the remainder of this section (i.e., Section A) provides
7 further details regarding the system-level aspects of the analysis strategy. Section B
8 provides additional details regarding the operation of the system. Section C discusses
9 exemplary applications of the system. And Section D describes an exemplary computing
10 environment for implementing features of the system.

11 To begin with, jumping ahead briefly to Fig. 2, this figure shows four exemplary
12 and non-limiting message-passing environments that can be investigated using the system
13 100 shown in Fig. 1. That is, the exemplary four message-passing environments shown
14 in Fig. 2 provide specific cases of the generic message-passing environment 102 shown
15 in Fig. 1.

16 Exemplary environment A (202) pertains to an intranet environment. In this
17 environment 202, a plurality of participants can communicate with each other via an
18 intranet 204. An intranet refers to a network that operates based on TCP/IP protocols
19 within the confines of an enterprise environment, such as a corporation or other
20 organization. A firewall prevents members outside of this environment from accessing
21 the resources of the intranet 204. The exemplary intranet 204 shown in Fig. 2 connects a
22 collection of client devices (e.g., clients 206, 208) with one or more servers (e.g., server
23 210). In this environment 202, the analysis system 108 can collect and analyze messages
24 transmitted between the clients (206, 208) or between the clients (206, 208) and the
25 server (210).

1 Exemplary environment B (212) pertains to a wide-area network environment. In
2 this environment 212, a plurality of participants can communicate with each other via the
3 Internet 214. The Internet refers to a network that operates based on TCP/IP protocols
4 and is accessible to a large number and generally unrestricted group of worldwide
5 participants. For purposes of illustration, the Internet 214 shown in Fig. 2 connects a
6 collection of client devices (e.g., clients 216, 218) with one or more servers (e.g., server
7 220). In this environment 212, the analysis system 108 can collect and analyze messages
8 transmitted between the clients (216, 218) or between the clients (216, 218) and the
9 server (220).

10 Environments A (202) and B (212) are not exhaustive of the network
11 environments that can be tested using the analysis system 108. Any kind of network
12 environment can be tested, including various LAN-type networks, Ethernet networks,
13 wireless networks, and so on. Further, the network environments (202, 212) shown in
14 Fig. 2 are highly simplified to facilitate discussion. In reality, these environments will
15 include other equipment, such as various routers, interfaces, gateways, and so on.

16 Exemplary environment C (222) pertains to a single machine including a plurality
17 of components, or one or more systems including a plurality of components. A
18 "component" as used herein can refer to any kind of equipment, such as a discrete data
19 processing device (e.g., a computer, memory device, router, etc.) or a part of a device
20 (such as a CPU, disk drive, RAM memory, various buses, external data stores, and so
21 on). In the simplified and illustrative case of Fig. 2, such a machine or a system includes
22 component A (224), component B (226), and component C (228) in cooperative
23 communication with each other via messages. Any one of these components can assume
24 the role of client, server, or some other role. In this environment 222, the analysis system
25 108 can collect and analyze messages transmitted between the components (224, 226, and

228). Such messages can thus be internal to the machine or the system. Accordingly, in this environment 222, collecting these messages may require access to the machine or system (and therefore the investigation of this environment 222 may be more intrusive compared to environments 202 and 212).

Exemplary environment D (230) pertains to a software module including a plurality of components. A “component” as used herein can refer to any collection of program instructions in any programming language, or any collection of declarative statements expressed in any declarative language (such as the extensible markup language, i.e., XML). In the simplified and illustrative case of Fig. 2, such a software module includes component A (232), component B (234), and component C (236) in cooperative communication with each other via messages, which may comprise functions calls, messages passed between objects in an object oriented language, and so on. Any one of these components can assume the role of client, server, or some other role. In this environment 230, the analysis system 108 can collect and analyze messages transmitted between the components (232, 234, and 236). Such messages can thus be internal to the machine or machines that implement the software program. Accordingly, like the last case 222, collecting these messages may require access to the machine(s).

Returning to the general depiction of the message-passing environment 102 in Fig. 1, the observation agents (110, 112, 114, 116) can be located throughout the environment 102. In one implementation, observation agents can be placed at locations that enable the analysis system 108 to intercept the messages between participants, e.g., after they are transmitted by a sender and before they are received by a receiver. In a network environment (such as environments 202 and 212), this can be performed by positioning the observation agents in the network at some intermediate point, such as a gateway, a router, at specialized monitoring equipment, or some other intermediary

1 location. This intermediary location can be associated with the sender entity, the
2 recipient entity, or some independent entity (such as the analyst). The entirety of the
3 transmitted messages can be captured or just parts of the messages (such as parts of the
4 headers or parts of the bodies of the messages).

5 In the machine environment (e.g., environment 222), messages can be intercepted
6 by monitoring information transmitted on lines coupling the components together, or
7 through some other mechanism.

8 In the code environment (e.g., environment 230), messages can be intercepted by
9 providing specialized software that extracts the messages during the execution of the
10 software, or through some other mechanism. For instance, this specialized software can
11 intercept messages passed to various subroutines, functions, software objects, interfaces,
12 buffers, logs, message stacks, etc.

13 In one implementation, the observation agents (110, 112, 114, 116) can be turned
14 on and off by a central administrator to suit different analysis needs. In this case, an
15 analyst can “turn off” those observation agents that are not needed, so as not to unduly
16 complicate the operation of the message-passing environment 102.

17 Whatever the case, Fig. 1 shows that each participant can include two observation
18 agents. A first observation agent can detect messages transmitted by the participant (as in
19 the case of observation agents 110 and 114), and a second observation agent can detect
20 messages received by the participant (as in the case of observation agents 112 and 116).
21 In other implementations, a single observation agent can be designed and/or positioned
22 within the network so as to record both inbound and outbound messages. In one case, the
23 observation agents (110, 112, 114, 116) can detect every message transmitted from or
24 received by the participants (104, 106) in a specified timeframe. In another case, the
25 observation agents (110, 112, 114, 116) may sample the messages transmitted from or

1 received by the participants (104, 106); the timing of this sampling can be governed by
2 predefined rules or can be random.

3 Messages can be transmitted to the data store 118 using any mechanism, e.g., via
4 hardwired and propriety communication lines, via any kind of network, via wireless
5 transmission, and so on.

6 The analysis system 108 itself can comprise any kind of data processing system,
7 such as a programmable computer device, a piece of equipment including hardwired
8 logic circuitry, or some combination of programmable computer and hardwired logic
9 circuitry. Generally, the analysis system 108 includes one or more processing units 122
10 (e.g., CPUs) and system memory 124 (e.g., Random Access Memory (RAM), etc.).
11 During operation, the memory 124 can store an operating system 126 that handles the
12 background tasks of the analysis system 108. The analysis system 126 can also store the
13 message analysis logic 120. The data store 118 can comprise any type of memory device
14 and any associated data management software associated therewith. The analysis system
15 108 may provide the data store 118 at a remote location with respect the message analysis
16 logic 120, or at the same location as the message analysis logic 120. The data store 118
17 itself can include a single repository of information or several distributed repositories of
18 information.

19 An analyst 128 interacts with the analysis system 108 via a collection of input
20 devices 130, such as a keyboard 132, mouse device 134, or other kinds of input device.
21 The analyst 128 also interacts with the analysis system 108 via display monitor 136.
22 Display monitor 136 can provide instructions to the analyst 128, receive input (e.g., via a
23 touch sensitive screen), and present analysis output results for reviewing by the analyst
24 128. The analysis system 108 can present the above-described information to the analyst
25 128 in the form of text output, a graphical user interface 138, or some other form. The

1 analysis system 108 can also output information to other devices, such as printers, remote
2 storage devices, remote computers, and so on.

3 Fig. 3 depicts the message analysis logic 120 and the data store 118 in greater
4 detail. The message analysis logic 120 can be implemented as a software program
5 comprising a plurality of program statements or declarative statements. This software
6 program, in turn, can be conceptualized as including a number of modules for handling
7 different functions performed by the message analysis logic 120. Each of these modules
8 can include a subset of the software program's instructions/statements.

9 Broadly speaking, message aggregation and conversion logic 302 receives
10 message information from the observation agents (110, 112, 114, 116) and aggregates
11 individual messages in this information into different groups. More specifically, a
12 message (M) can comprise a discrete chunk of information sent from a participant X to a
13 participant Y with a specific action (or command) and, optionally, other information. For
14 example, in network environments, a single message may be formatted using the Simple
15 Object Access Protocol (SOAP). SOAP provides a lightweight protocol to transfer
16 information over networks or other kind of distributed environments. This protocol
17 provides an extensible messaging framework using XML to provide messages that can be
18 sent on different kinds of underlying protocols. Each SOAP message includes a header
19 block and a body element. When transmitted over a network, the SOAP message may
20 also acquire additional header information attributed to protocols used by the network
21 (such as TCP/IP addressing information). Additional information regarding the SOAP
22 protocol is provided in the document *SOAP Version 1.2 Part 1: Messaging Framework*,
23 dated June 24, 2003, and available at W3C's web site. However, the transmission of
24 messages using SOAP is merely one illustrative example; other protocols and formats can
25 be used. Generally, in any format, a message can be conceptualized as including two

1 pieces of information: a first piece pertains to the transfer of information over the
2 exchange (such as message source, message destination, time, identification number(s),
3 etc.); and a second piece pertains to the specific operation or action being performed in
4 the message exchange (such as information regarding an online purchase, etc.). (The
5 action associated with the message can be gleaned from either the header or body of the
6 message.)

7 In one implementation, the message aggregation and conversion logic 302
8 receives message information from the participants in the form of "message traces." A
9 participant message trace refers to a series of messages originating from or sent to a
10 specific participant, ordered by time. For instance, participant 104 (shown in Fig. 1)
11 might send a trace to the message analysis logic 120 that contains ten minutes worth of
12 SOAP messages sent by it, and/or received by it. In one implementation, a trace may
13 contain all of the information in the intercepted messages. In another implementation, a
14 trace may contain only some information excerpted from the messages, such as
15 information extracted from the header and/or the body of SOAP messages. A trace may
16 or may not include an uninterrupted series of messages transmitted from or received by a
17 participant; for instance, in the case that information is collected from an observation
18 agent that only randomly samples messages, then the trace will not contain an
19 uninterrupted series of messages (that is, because some of the messages have not been
20 captured).

21 The traces are further arranged into so-called message sequences by the message
22 aggregation and conversion logic 302. The term "message sequence" is used liberally
23 herein to refer to any grouping of one or more messages received from the message-
24 passing environment 102 based on any criteria. For instance, a particular message
25 transaction between a client and server may require a series of messages between these

1 two participants. A message sequence can be compiled that corresponds to this sequence.
2 In another case, a message sequence can be compiled that pertains to messages
3 transmitted to or received by one or more participants in a specified time frame,
4 regardless of the nature of the transactions taking place. Still other bases for forming
5 sequences are possible based on other combinations of criteria. Generally, however, the
6 sequences are formed and ordered, at least in part, based on chronological information in
7 the messages.

8 More specifically, the operation of forming sequences may involve extracting
9 time information and/or other information from individual message traces, sorting the
10 messages based on such information, and grouping the messages into sequences based on
11 the results of the sorting. Additional information regarding this operation is provided in
12 the context of Fig. 4 (to be described below in turn).

13 The "conversion" component of the message aggregation and conversion logic
14 302 converts machine-specific identifying information associated with the messages into
15 logical or functional information associated with the respective roles that the machines
16 serve in the message-passing environment 102. For example, if a machine functions as a
17 client in a message exchange, then its machine-specific identifying code (that may be
18 present in the message sent or received by it) is converted to a functional identifier that
19 identifies this machine as a client. Additional information regarding this operation is also
20 provided below in the discussion of Fig. 4.

21 The output of the message aggregation and conversion logic 302 can be stored in
22 the data store 118. As shown in Fig. 3, the data store 118 includes a master collection
23 304 of message sequences, such as exemplary message sequence 306. As described
24 above, each message sequence can include one or more messages arranged by time
25 and/or other criteria.

1 Message sequence manager logic 308 generally manages the message sequence
2 information stored in the data store 118. This logic 308 can specifically cull specific
3 subsets of message sequences stored in the data store 118 based on specified criteria, and
4 then store these subsets in the data store 118 for subsequent analysis. For instance, the
5 data store 118 shows exemplary sequence subsets 310, 312 and 314. Subsets of
6 sequences can be formed based on time, transaction type, participants involved in the
7 message exchanges, and/or any other criteria depending on the objectives of the analyst
8 128 and the nature of the message-passing environment 102 involved.

9 Analysis logic 316 analyzes the one or more subsets of message sequences that
10 have been grouped together by the message sequence manager logic 308. The analysis
11 logic 316 can specifically perform cluster analysis on the sequences stored in the data
12 store 118 to group these sequences into different clusters based on specified criteria.
13 Alternatively, the analysis logic 316 can use other mechanisms for analyzing the
14 messages sequences, such as artificial intelligence analyses, neural network analyses,
15 various rule-based analyses, various kinds of statistical analyses, various kinds of pattern
16 matching analyses, and so on. Still alternatively, the analysis can be performed
17 manually, either in whole or in part, by a human analyst.

18 Finally, output logic 318 receives the results of the analysis logic 316 and
19 converts such output into an appropriate form for presentation to the analyst 128. For
20 instance, the output logic 318 can transform the output results for presentation in
21 graphical format, tabular format, or some other kind of format.

22 The operations performed in each of the above-described logic modules will be
23 described in greater detail in the next section.
24
25

B. Method of Operation

Fig. 4 illustrates an exemplary method 400 for performing message-based analysis using the system 100 of Fig. 1. In this figure, various algorithmic acts are summarized in individual “blocks.” Such blocks describe specific actions or decisions that are made or carried out as a process proceeds. Where a microcontroller (or equivalent) is employed, this method 400 provides a basis for a “control program” or software/firmware that may be used by such a microcontroller (or equivalent) to effectuate the desired control. In this case, the processes are implemented as machine-readable instructions or declarative statements storable in memory that, when executed by a processor, perform the various acts illustrated as blocks. While steps are shown as being performed in a prescribed order, it is possible to perform these steps in a different order.

Step 402: Collecting Traces

The method 400 begins in step 402, which entails collecting traces from participants in the message-passing environment 102. To arrange the messages based on time, it is necessary to associate time information with each captured message. In one case, time information is extracted from chronological information embedded in the messages themselves. This time might refer to when the message was created, when the message was sent, or based on some other information. Alternatively, or in addition, the observation agents (110, 112, 114, 116) can each provide a time stamp regarding when they intercepted the messages. Such time information may pertain to raw counter information, so it is useful to convert this information to more conventional time-based formats. Generally, because of the myriad of different ways that time can be extracted from the messages, it is necessary to arrive at a consistent methodology of interpreting time, and in turn, for synchronizing the different techniques for extracting time used in

1 the message-passing environment 102. It is also possible to capture and preserve time
2 information using multiple different techniques so as to provide multiple different
3 “views” of the behavior of the message-passing environment 102. Various heuristics can
4 also be used to assist in interpreting and harmonizing time information across traces; for
5 instance, a message is considered sent before it is received.

6 *Step 404: Converting to Logical Roles*

7 Step 404 entails converting the descriptive information that defines the
8 participants associated with the traces to more meaningful logical or functional
9 descriptions. For example, analysis system 108 may initially collect message traces that
10 identify the participants by machine-centric designators, such as “machine-012-xp” and
11 “machine-043-2k.” Step 404 converts these absolute descriptors into more functional
12 descriptors that describe the role that each participant serves in the transaction, such as
13 “client” or “server.” Such mapping of absolute descriptors to logical descriptors can be
14 performed by lookup mapping table, or user-assisted input. Alternatively, or in addition,
15 such mapping can be performed using automatic analysis of the traces to discover the role
16 that the participant is playing. For instance, such automatic analysis would classify a
17 participant that sends a request schedule message as a client because this behavior is
18 exhibited by a client and not a server.

19 Other logical designations besides client/server are possible. For instance, a peer-
20 to-peer network may not be structured using the client-server approach. In the general
21 case, the participants can be broken down into the broad category of sender and receiver;
22 however, even this does not hold true when a message is sent but never received by its
23 target. Further, in a broadcast/multicast mode of operation, a participant can send
24 messages to plural recipients.

25 *Steps 406 and 408: Forming Sequences*

1 Step 406 entails sorting the messages captured in the traces based on various
2 criteria, such as time, to form message sequences. The time synchronization provisions
3 discussed above are applied here to provide a consistent ordering of messages based on
4 time. Step 408 entails optionally storing the sequences formed in step 406 in a data store,
5 such as data store 118.

6 Steps 402-408 can be performed by the message aggregation and conversion logic
7 302 shown in Fig. 3, or in another module.

8 *Step 410: Grouping Sequences*

9 Step 410 entails selecting a group of sequences from the data store 118 for the
10 purpose of performing analysis on these sequences. For instance, the analyst 128 may be
11 primarily interested in investigating the behavior of a group of interacting participants at
12 a certain time of day. In this case, step 410 can cull a subset of sequences that provide
13 information regarding the participants of interest and the timeframe of interest.

14 Step 410 can be implemented using the message sequence manager logic 308
15 shown in Fig. 3.

16 *Step 412: Analyzing Sequences*

17 Step 412 entails actually performing analysis on the sequences selected in step
18 410. This step 412 can employ any type of analysis depending on the type of message-
19 passing environment 102 being analyzed, and depending on the objectives/interests of the
20 analyst 128. Exemplary types of analysis can include, but are not limited to: pattern
21 matching analyses; any kind of rule-based analyses; artificial intelligence analyses; any
22 kind of statistical analyses (such as cluster analysis); any type of neural network analyses,
23 and so forth.

24 To provide one exemplary example, step 412 will be described below in the
25 context of a cluster analysis strategy. Broadly stated, cluster analysis involves grouping

1 items in a set of items into one or more groups or clusters based on various criteria. Fig.
2 4 shows that the cluster analysis includes two broad steps: forming a data matrix (in step
3 414) and performing cluster analysis based on the thus formed data matrix (in step 416).
4 Each of these steps will be described below in greater detail.

5 As to step 414, a data matrix is formed from the selected message sequences to
6 emphasize different collections of information present in the message sequences. For
7 example, clustering can focus on specific re-try patterns, specific multi-response patterns,
8 specific transport fault conditions, specific gateway/firewall errors, etc. Generally, the
9 analyst 128 will typically select particular criteria for analysis based on the objectives of
10 the test and the characteristics of the subject message-passing environment 102. For
11 example, in one case, the analyst 128 may be interesting in performing functional tests to
12 discover whether there are "bugs" in a software program used by one or more of the
13 participants. In another case, the analyst 128 may be interested in investigating the
14 performance of the message-passing environment 102 in order to better tune such
15 environment 102 to improve its performance. Section C (below) provides additional
16 information regarding exemplary applications of the analysis techniques described herein.

17 Two exemplary techniques are discussed here for forming a data matrix on which
18 cluster analysis can be performed: feature-based techniques and similarity-based
19 techniques.

20 In feature-based techniques, step 414 takes each message sequence and extracts
21 numerical counts for different features present in the sequence. This could include
22 combinations of message command/action types (such as "Purchase" and "Sell" in web-
23 based commerce applications), sender/receiver pairs, properties of the message (e.g.,
24 "Secured" and "Reliable"), or application-level properties in the message (such as the
25 number of shares in financial-type applications). Action types can be extracted from

1 SOAP messages based on predefined XML information in the messages that specifies the
2 action types associated with the messages. Information regarding the action can also be
3 ascertained based on other parts of the messages, such as the HTTP header of the
4 message.

5 For instance, step 414 can extract features corresponding to counts of message
6 types. Consider, for example, the case of an illustrative sequence 0, in which a “request-
7 schedule” message has occurred ten times, while a “schedule-response” message has
8 occurred three times. The data matrix produced in this case would correspond to the
9 following:

10
11 **Exemplary Matrix Table 1**

12

Sequence	“request-schedule”	“schedule-response”	...
0	10	3	...

13
14
15

16 Another technique that can be used for extracting features involves counting
17 actions in pair-wise fashion between different participants in the message-passing
18 environment 102. An exemplary algorithm for implementing this technique is as follows:

19
20 **Exemplary Algorithm 1**

21
22 For each participant X:

23 For each participant Y:

24 For each action A:

25 Output From-To-A = “count A’s from X to Y”

In this algorithm, participants X and Y correspond to different messaging transmitting or receiving entities in the message-passing environment 102. However, in some cases, a participant X is the same as the participant Y, meaning that a single entity is both the transmitter and recipient of a message.

The following data matrix is produced using the above algorithm for exemplary participants labeled “C” and “S” (e.g., denoting client and server, respectively). The message actions appropriate to the exchange between these two participants are “request0” and “response0,” denoting a request made by one of the participants and a corresponding response made by the recipient of the request.

Exemplary Matrix Table 2

Seq- uence	C-C- request0	C-C- response0	C-S- request0	C-S- response0	S-C- request0	S-C- response0	S-S- request0	S-S- response0
0	0	0	10	0	0	3	0	0

In this message sequence, participant “C” made ten requests to participant “S” (as denoted by the column labeled “C-S request0.” (In other words, the notation “X-Y” indicates that the message action flows from entity “X” to entity “Y.”) Further, in this sequence, entity “S” responded to entity “C” three times, (as denoted by the column “S-C response0” column). Post processing can be performed to remove columns that do not list any actions (i.e., that list the number 0). As the reader will appreciate, in an actual message-passing environment, the number of columns produced using the pair-wise

1 approach described above may become relatively large. However, this does not
2 necessarily present an obstacle to efficient processing of such a matrix, as the processing
3 burden placed on some clustering algorithms grows, at worst, linearly with the number of
4 columns or dimensions.

5 Another exemplary approach is to perform logical time ordering of data stored in
6 the sequences. This approach can extract features depending on their chronological
7 occurrence in a specified timeframe. For example, this approach can extract information
8 depending on whether events took place before or after a specified point in time (denoted,
9 respectively, by the labels "happened-before" and "happened-after"). The following
10 algorithm constructs a data matrix based on such chronological considerations:

11 Exemplary Algorithm 2

12 For each participant X:

13 For each participant Y:

14 For each action A:

15 For each action B:

16 Output From-To-A-B-Before,

17 "count A's from X to Y which

18 happened-before B's from X to Y"

19 Output From-To-A-B-After,

20 "count A's from X to Y which

21 happened-after B's from X to Y"

22
23
24 This algorithm counts the number of actions "A" sent from participant X to participant Y
25 that happened before an action B is sent from participant X to participant Y. This

1 algorithm also counts the number of actions "A" sent from participant X to participant Y
2 that happened after an action B was sent from participant X to participant Y. For
3 instance, in the context of an online shopping message-passing environment, this
4 algorithm could be used to determine how many times that a user viewed a certain
5 category (or brand) of product before purchasing another category (or brand) of product.

6 Still another possible approach is to count the logical or physical time delays
7 between messages. The following algorithm extracts features based on a delay-based
8 paradigm:

10 Exemplary Algorithm 3

12 For each participant X:

13 For each participant Y:

14 For each action A:

15 For each header H:

16 Output From-To-A-H, "count A's
17 containing H from X to Y"

18
19 This algorithm counts action A's sent from X to Y providing that they have certain
20 parameters in their header H (or fall within a certain range of such parameters). For
21 instance, time information can be extracted from an IP header, SOAP header, or other
22 kind of message network header. Alternatively, time information can be inferred from
23 the time that the message was intercepted, as determined by the observation agent. Still
24 other techniques are available for gauging time information from messages. Using this
25

1 chronological information, it is possible to determine how long certain actions take to
2 perform, or the amount of time between different actions, and so forth.

3 Other algorithms can be devised to extract different features from the messages
4 depending on the objectives of the analyst 128, the type of message-passing environment
5 102 involved, the composition of messages, and/or other factors. In any event, the output
6 of such feature extraction constitutes a multi-dimensional data matrix. Clusters are
7 formed based on information in this matrix, as will be discussed in the context of step
8 416.

9 Still referring to step 414, the second technique for forming a data matrix is
10 similarity-based analysis of the messages. In this technique, instead of directly extracting
11 features from the sequences, each sequence is compared with other sequences to derive
12 difference values that express differences between information associated with the
13 sequences. That is, assume that messages X and Y include parameters x1 and y1,
14 respectively. A data matrix is computed using the similarity technique by subtracting x1
15 from y1 to derive a difference value d. The algorithm can normalize the difference value
16 by defining the similarity as: $\text{similarity} = \text{MaximumValue} / (\text{Calculated_Difference}(x, y)$
17 $+ 1.0)$, where the Calculated_Difference variable should return a value d such that $0 \leq d$
18 $\leq \text{MaximumValue}$.

19 A variety of difference algorithms can be applied to calculate a similarity matrix,
20 such as string/sequence matching. In this approach, if a message was not sent, the
21 algorithm increases the difference count by M, and if a message was sent twice, the
22 algorithm increases the difference count by N, and so on.

23 With the similarity technique, it is also possible to compare a set of sequences
24 with a known sequence that has been collected and stored in advance. This known
25 sequence may represent a baseline sequence that the analyst 128 is confident represents

1 the proper or optimal functioning of the message-passing environment 102. In this case,
2 the analyst 128 can form a difference matrix that reflects the deviation of the message-
3 passing environment 102 being tested from the baseline known sequence. For example,
4 using this technique, the analyst 128 can compare a "good" server trace with a
5 measured/observed trace, or a known "bad" server trace with a measured/observed trace.
6 In the former case, a sequence that diverges from a good server trace cluster might be
7 indicative of a failed server; in the latter case, a sequence that is grouped with the bad
8 server trace cluster might be indicated of a failed server.

9 In another case, the known sequence can be collected from another kind of
10 message-passing environment, such as a related type of message-passing environment.
11 In this scenario, an analyst can form a difference matrix that reflects how the message-
12 passing environment 102 under consideration differs from related systems, such as
13 systems produced by different computer or software manufacturers, or systems
14 employing different processing strategies or software application versions. Such system-
15 to-system comparisons may be particularly useful in analyzing specific re-try patterns,
16 specific multi-response patterns, specific transport fault conditions, specific
17 gateway/firewall errors, and so on. For example, an analyst can use this comparison
18 technique to compare the behavior of two software programs (e.g., a Stock Purchase
19 program and a Calendar program) that run on the same network configuration, even
20 though the messages propagated between participants in these environments have
21 different application-related content.

22 Once the data matrix has been formed, step 416 comes into play by forming
23 clusters on the basis of information in the data matrix. Any type of clustering algorithm
24 can be used to perform this task, such as algorithms using the partitional paradigm,
25 agglomerative paradigm, graph-partitioned paradigm, etc. For example, one suite of

1 clustering strategies that can be used is provided the CLUTO software package provided
2 by George Karypis (Department of Computer Science & Engineering, Twin Cities
3 Campus, University of Minnesota, Minneapolis, Minnesota), which employs all of the
4 above paradigms. The clustering step 416 can rely on one clustering algorithm to analyze
5 the data set, or can combine several different clustering algorithms. In the latter case, the
6 algorithm can automatically select the best approach by trying each one, or can combine
7 the results of different approaches, or can iteratively converge on an optimal solution by
8 repeating the clustering analysis with different settings or approaches.

9 In any case, the analyst 128 can control the clustering algorithm by selecting the
10 number of clusters that should be created. In one implementation, the analyst 128 may
11 want the clustering algorithm to group the sequences into clusters such that the ratio of
12 the number of clusters produced to the number of initial sequences is about 15%. That is,
13 if 100 sequences are used to form the data matrix, then the algorithm should produce
14 about 15 clusters that group these sequences together.

15 Other settings allow the analyst 128 to specify the techniques used by the
16 clustering algorithm to measure distances between clustered objects or the distances
17 between objects and the clusters to which they are associated. For example, the analyst
18 128 may specify that the algorithm should compute this distance based on the square root
19 of the distance between two objects instead of a normal distance. Alternatively, the
20 analyst 128 may specify that the algorithm should measure the distance from an object to
21 the nearest neighboring object in the cluster, or measure the distance from the farthest
22 neighboring object in the cluster, or measure the distance from the weighted center of the
23 cluster, and so on.
24
25

1 The output of step 416 comprises a listing of clusters and the sequences
2 associated therewith. For instance, consider the case where seven sequences (numbered 0
3 through 7) were fed to the clustering algorithm. In this case, the output might be:

4
5 Cluster 0: Sequence 0, 5, 6

6 Cluster 1: Sequence 1, 2, 4

7 Cluster 2: Sequence 3
8

9 The above seven sequences might contain known reference sequences added to
10 the group of sequences to assist in interpreting the results. Known reference sequences
11 can correspond to sequences that reflect the error-free operation of the message-passing
12 environment 102, or known failure conditions within the environment 102.

13 To repeat, step 412 is not limited to cluster analysis; other techniques can be used.
14 For example, step 412 can compare the message sequences against a formal model of the
15 system (e.g., provided by a state machine). This comparison can place each sequence in
16 one of two "clusters," corresponding respectively to whether each sequence adheres to
17 the model or does not adhere to the model.

18 *Step 418: Post-Analyzing and/or Presenting Results*

19 Step 418 involves optionally performing additional analysis on the output of step
20 412. In the event clustering analysis was used in step 412, step 418 may entail
21 performing post-analysis to select sequences that are "interesting." Generally, the term
22 "interesting" means different things depending on the objectives of the analyst 128. The
23 analyst 128 might consider a sequence interesting because it is suggestive of a functional
24 or performance-related error. Alternatively, the analyst 128 may be interested in
25 identifying message sequences that are indicative of beneficial phenomena, such as

1 instances when a message-passing environment performs particularly well. Still
2 alternatively, the analyst 128 may be interested in identifying trends in activity within the
3 environment for strictly marketing-related purposes. Section C below provides additional
4 examples of possible applications of the method 400 shown in Fig. 4.

5 Whatever the analyst 128's objectives, the post-processing can entail a variety of
6 techniques. The techniques can use automatic analysis of formed clusters using various
7 rule-based systems, artificial intelligence systems, neural network systems, and so forth.
8 Alternatively, the techniques can provide a visual presentation of the clusters to the
9 analyst 128 and allow the analyst 128 to manually select interesting sequences based on
10 his or her own informed judgment. Still alternatively, the post-analysis can comprise a
11 combination of automated and manual techniques.

12 For example, step 418 may sort the formed clusters on the basis of the number of
13 members in the clusters (from smallest to largest). The analyst 128 may then want to
14 further examine the first N % of clusters in this ranked list. This is because small clusters
15 of sequences may be indicative of particularly anomalous or interesting conditions that
16 warrant further investigation. Clusters with only one member (i.e., singleton clusters)
17 tend to be especially interesting. A small cluster does not necessarily represent an error
18 or performance problem; however, such a small cluster has at least some feature or
19 features which make it stand out from the other clusters.

20 Fig. 5 shows an exemplary output of the method 400 of Fig. 4. The output
21 consists of a two-dimensional presentation of the formed clusters (502-510). The axes of
22 the graph can correspond to different attributes of the sequences. However, in other
23 cases, the method 400 can present the output of the clustering process in another format,
24 such as a table that simply ranks the clusters based on number of members in the clusters.
25

1 In the illustrative case shown in Fig. 5, clusters 502, 506, and 508 contain a
2 relatively large number of members, while clusters 504 and 510 contain relatively few
3 members. Hence, the analyst 128 might be particularly interested in performing further
4 analysis on the sequences contained in clusters 504 and 510. The system 100 shown in
5 Fig. 1 can partially automate this further analysis by linking each cluster to information
6 regarding the sequences associated with the cluster. This can be performed via hypertext
7 links or some other linking mechanism. More specifically, the system 100 could provide
8 supplemental information such as information listing the actual messages in the identified
9 sequences. Additionally, the system 100 could be configured to perform additional
10 automated analysis on the selected clusters upon the request of the analyst 128.

11 Various graphical aids could also be provided. For instance, the system 100 can
12 present a schematic of the message-passing environment 102. Mapping logic can be
13 provided that correlates interesting sequences with locations in the schematic
14 corresponding to agents (participants) that may be associated with the interesting
15 sequences. This might be particularly useful in identifying equipment that may be
16 performing incorrectly or poorly.

17 18 **C. Exemplary Applications**

19 The analyst 128 can apply the method 400 shown in Fig. 4 to a great variety of
20 investigative tasks. In one case, the analyst 128 might be interested in identifying
21 sequences that either represent functional errors (e.g., the environment is producing
22 inaccurate results), or performance-related problems (e.g., the environment may be
23 producing accurate results, but is producing them in a substandard manner, that is, either
24 too slow or by consuming too much memory, etc.).
25

1 Consider, for example, the following sequences produced by an environment that
2 involves performing arithmetic operations (e.g., using a well-known GUI-based
3 calculator program). The client and server mentioned below might refer to separate
4 computers coupled together via a network, or separate modules within a single computer.

5 Sequence 1: Client sends message ("add 1, 2"), and server sends response
6 ("3").

7 Sequence 2: Client sends message ("add 3, 4") but must retry sending ten
8 times. Server is too busy to respond to the first nine requests, but finally sends
9 one response to the tenth request ("7").

10 Sequence 3: Client sends message ("plus 1, 2"), and server sends failure
11 ("not supported").

12 Sequence 4: Client sends message ("plus 3, 8"), but server is too busy and
13 sends no response.

14 In these executions, the analyst 128 might be particularly interested in further
15 examining sequences 2 and 4. This is because these cases have fundamentally different
16 message exchange patterns compared to cases 1 and 3. Anomalous conditions might
17 become even clearer upon collecting and analyzing a larger population of sequences.
18 Generally, the method 400 can be used to identify outright coding errors, or to identify
19 lack of coding sophistication (such as poor handling of re-try logic). The results can be
20 used for debugging, for improving algorithms, and for deploying new policies that govern
21 the message-passing environment 102.

22 The method 400 can also be used to identify transient circumstances that affect
23 behavior yet may not be attributable to the participants that originate or receive the
24 messages. For instance, consider the case where participant X sends a message to
25 recipient Y through two different interface routes. One of these routes might perform

1 substantially worse than the other. The method 400 can provide information which
2 assists the analyst 128 in pinpointing the equipment that may be responsible for this
3 discrepancy. For instance, the analyst 128 may come to the conclusion that a gateway is
4 involved in one route that is performing poorly, e.g., by dropping packets. Such a
5 conclusion can be reached even though the gateway may not affect the content of the
6 messages being transmitted.

7 In still another application, the analyst 128 may be interested in identifying cases
8 in which the environment performs particular well. The analyst 128 might want to study
9 this phenomenon to determine what contributes to its success, so that this condition
10 attributed to success can be duplicated in other parts of the environment on a more
11 consistent basis.

12 Another application is to detect anomalous conditions in the message-passing
13 environment 102 that may be suggestive of improper use of the environment. For
14 example, the method 400 can be used to detect patterns of message exchange that are
15 indicative of unauthorized access to network resources or fraudulent activity. Such
16 patterns can emerge by investigating outlying clusters or small clusters. Also, the analyst
17 128 can interject known message patterns that are indicative of improper conduct into the
18 analysis. In this case, the method 400 can provide an indication of improper conduct if it
19 classifies collected message sequences with known "bad" sequences.

20 More generally, in this domain of analysis, the sequence of received messages is
21 often as significant for analysis as the number of messages. The firewall used in a
22 network environment might not be able to filter out prohibited message patterns because
23 it operates using a stateless paradigm, and therefore is incapable of recognizing the
24 connection between messages. Consider the case where the firewall may permit the
25 exchange of both create-dialog and teardown chat-session messages, but a message

1 sequence consisting of 10,000 teardown chat-session messages, one create-dialog
2 message, and 10,000 more teardown chat-session messages might be suggestive of
3 improper activity; being stateless, the firewall might not be able to detect this problem, but
4 the above-described method 400 can pick out this pattern.

5 Another application of the method is in the field of marketing. For instance, the
6 analyst 128 may be primarily concerned with the patterns of purchasing behavior
7 exhibited by users, rather than whether the message-passing environment is working
8 properly. For instance, an analyst 128 can use the method 400 to determine various
9 correlations relating to users' web browsing activities or online shopping activities. The
10 method 400 can determine whether certain activities are prevalent in certain time periods,
11 whether certain activities are associated with the other activities or events, and so on.
12 The analyst 128 could use this information to improve the dissemination of products and
13 services to individuals assessed to be most likely desirous of purchasing such products
14 and services. The method 400 also provides a mechanism for non-commercial research
15 (such as various academic or government-related studies of web usage).

16 The above applications are not limitative of the many uses of the method 400
17 shown in Fig. 4.

18 The benefits of this approach are likewise diverse. As explained above, one
19 advantage is that the analyst 128 need not gain access to the equipment and systems
20 being tested in order to analyze them. (However, the analyst 128 may have to take a
21 more intrusive approach when analyzing the messages passed between components in a
22 single machine, or between modules of program code; this is because these message
23 events might not be accessible "on a wire" to parties that do not have direct access to the
24 machine or program under investigation.)
25

D. Exemplary Computer Environment

Fig. 6 provides additional information regarding a computer environment 600 that can be used to implement the analysis system 108 shown in Fig. 1. That is, the computing environment 600 includes the general purpose computer 108 and the display device 136 discussed in the context of Fig. 1. However, the computing environment 600 can include other kinds of computer and network architectures. For example, although not shown, the computer environment 600 can include hand-held or laptop devices, set top boxes, programmable consumer electronics, mainframe computers, gaming consoles, etc. Further, Fig. 6 shows elements of the computer environment 600 grouped together to facilitate discussion. However, the computing environment 600 can employ a distributed processing configuration. In a distributed computing environment, computing resources can be physically dispersed throughout the environment.

Exemplary computer 108 includes one or more processors or processing units 122, a system memory 124, and a bus 602. The bus 602 connects various system components together. For instance, the bus 602 connects the processor 122 to the system memory 124. The bus 602 can be implemented using any kind of bus structure or combination of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. For example, such architectures can include an Industry Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA) local bus, and a Peripheral Component Interconnects (PCI) bus also known as a Mezzanine bus.

Computer 108 can also include a variety of computer readable media, including a variety of types of volatile and non-volatile media, each of which can be removable or non-removable. For example, system memory 124 includes computer readable media in

1 the form of volatile memory, such as random access memory (RAM) 604, and non-
2 volatile memory, such as read only memory (ROM) 606. ROM 606 includes an
3 input/output system (BIOS) 608 that contains the basic routines that help to transfer
4 information between elements within computer 108, such as during start-up. RAM 604
5 typically contains data and/or program modules in a form that can be quickly accessed by
6 processing unit 122.

7 Other kinds of computer storage media include a hard disk drive 610 for reading
8 from and writing to a non-removable, non-volatile magnetic media, a magnetic disk drive
9 612 for reading from and writing to a removable, non-volatile magnetic disk 614 (e.g., a
10 "floppy disk"), and an optical disk drive 616 for reading from and/or writing to a
11 removable, non-volatile optical disk 618 such as a CD-ROM, DVD-ROM, or other
12 optical media. The hard disk drive 610, magnetic disk drive 612, and optical disk drive
13 616 are each connected to the system bus 602 by one or more data media interfaces 620.
14 Alternatively, the hard disk drive 610, magnetic disk drive 612, and optical disk drive 616
15 can be connected to the system bus 602 by a SCSI interface (not shown), or other
16 coupling mechanism. Although not shown, the computer 108 can include other types of
17 computer readable media, such as magnetic cassettes or other magnetic storage devices,
18 flash memory cards, CD-ROM, digital versatile disks (DVD) or other optical storage,
19 electrically erasable programmable read-only memory (EEPROM), etc.

20 Generally, the above-identified computer readable media provide non-volatile
21 storage of computer readable instructions, data structures, program modules, and other
22 data for use by computer 108. For instance, the readable media can store the operating
23 system 126, one or more application programs 622 (such as the message analysis logic
24 120), other program modules 624, and program data 626.

1 The computer environment 600 can include a variety of input devices. For
2 instance, the computer environment 600 includes the keyboard 132 and a pointing device
3 134 (e.g., a “mouse”) for entering commands and information into computer 108. The
4 computer environment 600 can include other input devices (not illustrated), such as a
5 microphone, joystick, game pad, satellite dish, serial port, scanner, card reading devices,
6 digital or video camera, etc. Input/output interfaces 628 couple the input devices to the
7 processing unit 122. More generally, input devices can be coupled to the computer 108
8 through any kind of interface and bus structures, such as a parallel port, serial port, game
9 port, universal serial bus (USB) port, etc.

10 The computer environment 600 also includes the display device 136. A video
11 adapter 630 couples the display device 136 to the bus 602. In addition to the display
12 device 136, the computer environment 600 can include other output peripheral devices,
13 such as speakers (not shown), a printer (not shown), etc.

14 Computer 108 can operate in a networked environment using logical connections
15 to one or more remote computers, such as a remote computing device 632. The remote
16 computing device 632 can comprise any kind of computer equipment, including a general
17 purpose personal computer, portable computer, a server, a router, a network computer, a
18 peer device or other common network node, etc. Remote computing device 632 can
19 include all of the features discussed above with respect to computer 108, or some subset
20 thereof.

21 Any type of network can be used to couple the computer 108 with remote
22 computing device 632, such as a local area network (LAN) 634, or a wide area network
23 (WAN) 636 (such as the Internet). When implemented in a LAN networking
24 environment, the computer 108 connects to local network 634 via a network interface or
25 adapter 638. When implemented in a WAN networking environment, the computer 108

1 can connect to the WAN 636 via a modem 640 or other connection strategy. The modem
2 640 can be located internal or external to computer 108, and can be connected to the bus
3 602 via serial I/O interfaces 642 other appropriate coupling mechanism. Although not
4 illustrated, the computing environment 600 can provide wireless communication
5 functionality for connecting computer 108 with remote computing device 632 (e.g., via
6 modulated radio signals, modulated infrared signals, etc.).

7 In a networked environment, the computer 108 can draw from program modules
8 stored in a remote memory storage device 644. Generally, the depiction of program
9 modules as discrete blocks in Fig. 6 serves only to facilitate discussion; in actuality, the
10 programs modules can be distributed over the computing environment 600, and this
11 distribution can change in a dynamic fashion as the modules are executed by the
12 processing unit 904.

13 Wherever physically stored, one or more memory modules 124, 614, 618, 644,
14 etc. can be provided to store the message analysis logic 120 shown in Figs. 1 and 3.
15

16 Although the invention has been described in language specific to structural
17 features and/or methodological acts, it is to be understood that the invention defined in
18 the appended claims is not necessarily limited to the specific features or acts described.
19 Rather, the specific features and acts are disclosed as exemplary forms of implementing
20 the claimed invention.
21
22
23
24
25